# Distributed Learning with Infinitely Many Hypotheses

Angelia Nedić, Alex Olshevsky and César A. Uribe

arXiv:1605.02105v1 [math.OC] 6 May 2016

**Abstract**

We consider a distributed learning setup where a network of agents sequentially access realizations of a set of random variables with unknown distributions. The network objective is to find a parametrized distribution that best describes their joint observations in the sense of the Kullback-Leibler divergence. Apart from recent efforts in the literature, we analyze the case of countably many hypotheses and the case of a continuum of hypotheses. We provide non-asymptotic bounds for the concentration rate of the agents' beliefs around the correct hypothesis in terms of the number of agents, the network parameters, and the learning abilities of the agents. Additionally, we provide a novel motivation for a general set of distributed Non-Bayesian update rules as instances of the distributed stochastic mirror descent algorithm.

## I. Introduction

Sensor networks have attracted massive attention in past years due to its extended range of applications and its ability to handle distributed sensing and processing for systems with inherently distributed sources of information, e.g., power networks, social, ecological and economic systems, surveillance, disaster management health monitoring, etc. [1], [2], [3]. For such distributed systems, one can assume complete communication between every source of information (e.g. nodes or local processing unit) and centralized processor can be cumbersome. Therefore, one might consider cooperation strategies where nodes with limited sensing capabilities distributively aggregate information to perform certain global estimation or processing task.

The authors are with the Coordinated Science Laboratory, University of Illinois, 1308 West Main Street, Urbana, IL 61801, USA, {angelia,aolshev2,cauribe2}@illinois.edu. This research is supported partially by the National Science Foundation under grants no. CCF 11-11342 and no. CMMI-1463262 and by the Office of Naval Research under grant no. N00014-12-1-0998.

Following the seminal work of Jadbabaie et al. in [4], [5], there have been many studies of Non-Bayesian rules for distributed algorithms. Non-Bayesian algorithms involve an aggregation step, usually consisting of weighted geometric or arithmetic average of the received beliefs, and a Bayesian update that is based on the locally available data. Therefore, one can exploit results from consensus literature [6], [7], [8], [9], [10] and Bayesian learning literature [11], [12]. Recent studies have proposed several variations of the Non-Bayesian approach and have proved consistent, geometric and non-asymptotic convergence rates for a general class of distributed algorithms; from asymptotic analysis [13], [14], [15], [16], [17], [18] to non-asymptotic bounds [19], [20], [21], time-varying directed graphs [22] and transmission and node failures [23].

In contrast with the existing results that assume a finite hypothesis set, in this paper, we are extending the framework to the cases of a countable many and a continuum of hypotheses. We build upon the work in [24] on non-asymptotic behaviors of Bayesian estimators to construct non-asymptotic concentration results for distributed learning. In the distributed case, the observations will be scattered among a set of nodes or agents and the learning algorithm should guarantee that every node in the network will learn the correct parameter as if it had access to the complete data set. Our results show that in general the network structure will induce a transient time after which all agents will learn at a network independent rate, where the rate is geometric.

*The contributions of this paper are threefold*: First, we provide an interpretation of a general class of distributed Non-Bayesian algorithms as specific instances of a distributed version of the stochastic mirror descent. This motivates the proposed update rules and makes a connection between the Non-Bayesian learning literature in social networks and the Stochastic Approximations literature. Second, we establish a non-asymptotic concentration result for the proposed learning algorithm when the set of hypothesis is countably infinite. Finally, we provide a non-asymptotic bound for the algorithm when the hypothesis set is a bounded subset of $\mathbb{R}^d$. This is an initial approach to the analysis of distributed Non-Bayesian algorithms for a more general family of hypothesis sets.

This paper is organized as follows: Section II describes the studied problem and the proposed algorithm, together with the motivation behind the proposed update rule and its connections with distributed stochastic mirror descent algorithm. Section IV and Section V provide the non-asymptotic concentration rate results for the beliefs around the correct hypothesis set for the cases of countably many and continuum of hypotheses, respectively. Finally, conclusions are

presented in Section VI.

**Notation**: The set $B^c$ denotes the complement of a set $B$. Notation $\mathbb{P}_B$ and $\mathbb{E}_B$ denotes the probability measure and expectation under a distribution $P_B$. The $ij$-th entry of a matrix $A$ is denoted by $[A]_{ij}$ or $a_{ij}$. Random variables are denoted with upper-case letters, while the corresponding lower-case letters denote their realizations. Time indices are indicated by subscripts and the letter $k$. Superscripts represent the agent indices, which are usually $i$ or $j$.

## II. Problem Formulation

We consider the problem of distributed non-Bayesian learning, where a network of agents access sequences of realizations of a random variable with an unknown distribution. The random variable is assumed to be of finite dimension with the constraint that each agent can access only a strict subset of the entries of the realizations (e.g., an $n$-dimensional vector and $n$ agents each observing a single entry). Observations are assumed to be independent among the agents. We are interested in situations where no single agent has the ability to learn the underlying distribution from its own observations, while collectively the agents can do so if they collaborate. The learning objective is for the agents to jointly agree on a distribution (from a parametrized family of distributions or a hypothesis set) that best describes the observations in a specific sense (e.g., Kullback-Leibler divergence). Therefore, the distributed learning objective requires collaboration among the agents which can be ensured by using some protocols for information aggregation and coordination. Specifically in our case, agent coordination consists of sharing their estimates (*beliefs*) of the best probability distribution over the hypothesis set.

Consider, for example, the distributed source location problem with limited sensing capabilities [25], [26]. In this scenario a network of $n$ agents receives noisy measurements of the distance to a source, where sensing capabilities of each sensor might be limited to a certain region. The group objective is to jointly identify the location of the source and that every node knows the source location. Figure 1 shows an example, where a group of 7 agents (circles) wants to localize a source (star). There is an underlying graph that indicates the communication abilities among the nodes. Moreover, each node has a sensing region indicated by the dashed line around it. Each agent $i$ obtains realizations of the random variable $S_k^i = \|x^i - \theta^*\| + W_k^i$, where $\theta^*$ is the location of the source, $x^i$ is the position of agent $i$ and $W_k^i$ is a noise in the observations. If we consider $\Theta$ as the set of all possible locations of the source, then each $\theta \in \Theta$ will induce a

probability distribution about the observations of each agent. Therefore, agents need to cooperate and share information in order to guarantee that all of them correctly localize the target.

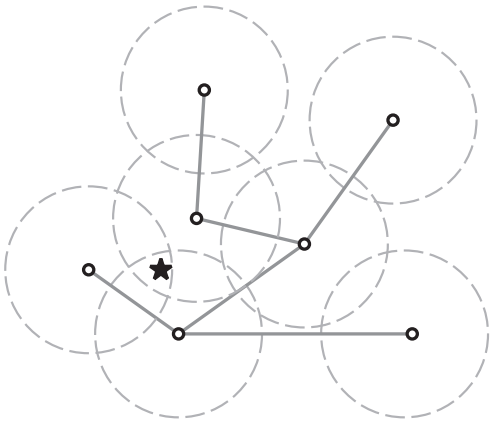

Fig. 1: Distributed source localization example

We will consider a more general learning problem, where agent observations are drawn from an unknown joint distribution $\boldsymbol{f} = \prod_{i=1}^{n} f^i$, where $f^i$ is the distribution governing the observations of agent $i$. We assume that $\boldsymbol{f}$ is an element of $\mathscr{P} = \prod_{i=1}^{n} \mathscr{P}^i$, the space of all joint probability measures for a set of $n$ independent random variables $\{S^i\}_{i=1}^{n}$ (i.e., $S^i$ is distributed according to an unknown distribution $f^i$). Also, we assume that each $S^i$ takes values in a finite set. When these random variables are considered at time $k$, we denote them by $S_k^i$.

Later on, for the case of countably many hypotheses, we will use the *pre-metric* space $(\mathscr{P}, D_{KL})$ as the vector space $\mathscr{P}$ equipped with the Kullback-Liebler divergence. This will generate a topology, where we can define an open ball $\mathcal{B}_r(p)$ with a radius $r > 0$ centered at a point $p \in \mathscr{P}$ by $\mathcal{B}_r(p) = \{q \in \mathscr{P} | D_{KL}(q, p) < r\}$. When the set of hypothesis is continuous, we instead equip $\mathscr{P}$ with the Hellinger distance $h$ to obtain the *metric* space $(\mathscr{P}, h)$, which we use to construct a special covering of subsets $B \subset \mathscr{P}$ consisting of $\delta$-separated sets.

Each agent constructs a set of hypothesis parametrized by $\theta \in \Theta$ about the distribution $f^i$. Let $\mathscr{L}^i = \{P_\theta^i | \theta \in \Theta\}$ be a parametrized family of probability measures for $S_k^i$ with densities $\ell^i(\cdot|\theta) = dP_\theta^i / d\lambda^i$ with respect to a dominating measure[1] $\lambda^i$. Therefore, the learning goal is to

[1]A measure $\mu$ is dominated by $\lambda$ if $\lambda(B) = 0$ implies $\mu(B) = 0$ for every measurable set $B$.

distributively solve the following problem:

$$\min_{\theta\in\Theta} F(\theta) \triangleq D_{KL}\left(\boldsymbol{f}\|\boldsymbol{\ell}\left(\cdot|\theta\right)\right) \tag{1}$$
$$= \sum_{i=1}^{n} D_{KL}\left(f^i\|\ell^i\left(\cdot|\theta\right)\right)$$

where $\boldsymbol{\ell}\left(\cdot|\theta\right) = \prod_{i=1}^{n}\ell^i\left(\cdot|\theta\right)$ and $D_{KL}\left(f^i\|\ell^i\left(\cdot|\theta\right)\right)$ is the Kullback-Leibler divergence between the true distribution of $S_k^i$ and $\ell^i(\cdot|\theta)$ that would have been seen by agents $i$ if hypothesis $\theta$ were correct. For simplicity *we will assume that there exists a single* $\theta^* \in \Theta$ such that $\ell^i\left(\cdot|\theta^*\right) = f^i$ almost everywhere for all agents. Results readily extends to the case when the assumption does not hold (see, for example, [27], [20], [22] which disregard this assumption).

The problem in Eq. (1) consists of finding the parameter $\theta^*$ such that $\boldsymbol{\ell}\left(\cdot|\theta\right) = \prod_{i=1}^{n}\ell^i\left(\cdot|\theta\right)$ minimizes its Kullback-Liebler divergence to $\boldsymbol{f}$. However, $\mathscr{L}^i$ is only available to agent $i$ and the distribution $\boldsymbol{f}$ is unknown. Agent $i$ gains information on $f^i$ by observing realizations $s_k^i$ of $S_k^i$ at every time step $k$. The agent uses these observations to construct a sequence $\{\mu_k^i\}$ of probability distributions over the parameter space $\Theta$. We refer to these distributions as agent $i$ *beliefs*, where $\mu_k^i\left(B\right)$ denotes the belief, at time $k$, that agent $i$ has about the event $\theta^* \in B \subseteq \Theta$ for a measurable set $B$.

We make use of the following assumption.

**Assumption 1** *For all agents* $i = 1, \ldots, n$ *we have:*

*(a) There is a unique hypothesis* $\theta^*$ *such that* $\ell^i\left(\cdot|\theta^*\right) = f^i$.

*(b) If* $f^i\left(s^i\right) > 0$*, then there exists an* $\alpha > 0$ *such that* $\ell^i\left(s^i|\theta\right) > \alpha$ *for all* $\theta \in \Theta$.

Assumption 1(a) guarantees that we are working on the realizable case and there are no conflicting models among the agents, see [27], [20], [22] for ways of how to remove this assumption. Moreover in Assumption 1(b), the lower bound $\alpha$ assumes the set of hypothesis are dominated by $f^i$ (i.e., our hypothesis set is absolutely continuous with respect to the true distribution of the data) and provide a way to show bounded differences when applying the concentration inequality results.

Agents are connected in a network $\mathcal{G} = \{V, E\}$ where $V = \{1, 2, \ldots, n\}$ is the set of agents and $E$ is a set of undirected edges, where $(i, j) \in E$ if agents $i$ and $j$ can communicate with each other. If two agents are connected they share their beliefs over the hypothesis set at every time

instant $k$. We will propose a distributed protocol to define how the agents update their beliefs based on their local observations and the beliefs received from their neighbors. Additionally, each agent weights its own belief and the beliefs of its neighbors; we will use $a_{ij}$ to denote the weight that agent $i$ assigns to beliefs coming from its neighbor $j$, and $a_{ii}$ to denote the weight that the agent assigns to its own beliefs. The assumption of static undirected links in the network is made for simplicity of the exposition. The extensions of the proposed protocol to more general cases of time varying undirected and directed graphs can be done similar to the work in [27], [20], [22].

Next we present the set of assumptions on the network over which the agents are interacting.

**Assumption 2** *The graph $\mathcal{G}$ and matrix $A$ are such that:*

*(a) $A$ is doubly-stochastic with $[A]_{ij} = a_{ij} > 0$ if $(i, j) \in E$.*

*(b) If $(i, j) \notin E$ for some $i \neq j$ then $a_{ij} = 0$.*

*(c) $A$ has positive diagonal entries, $a_{ii} > 0$ for all $i \in V$.*

*(d) If $a_{ij} > 0$, then $a_{ij} \geq \eta$ for some positive constant $\eta$.*

*(e) The graph $\mathcal{G}$ is connected.*

Assumption 2 is common in distributed optimization and consensus literature. It guarantees convergence of the associated Markov Chain and defines bounds on relevant eigenvalues in terms of the number $n$ of agents. To construct a set of weights satisfying Assumptions 2, for example, one can consider a lazy metropolis (stochastic) matrix of the form $\bar{A} = \frac{1}{2}I + \frac{1}{2}A$, where $I$ is the identity matrix and $A$ is a stochastic matrix whose off-diagonal entries satisfy

$$a_{ij} = \begin{cases} \frac{1}{\max\{d^i+1, d^j+1\}} & \text{if } (i,j) \in E \\ 0 & \text{if } (i,j) \notin E \end{cases}$$

where $d^i$ is the degree (the number of neighbors) of node $i$. Generalizations of Assumption 2 to time-varying undirected are readily available for weighted averaging and push-sum approaches [28], [10], [9].

## III. Distributed Learning Algorithm

In this section, we present the proposed learning algorithm and a novel connection between Bayesian update and the stochastic mirror descent method. We propose the following (theoretical)

algorithm, where each node updates its beliefs on a measurable subset $B \subseteq \Theta$ according to the following update rule: for all agents $i$ and all $k \geq 1$,

$$\mu_k^i(B) = \frac{1}{Z_k^i} \int_{\theta \in B} \prod_{j=1}^{n} \left( \frac{d\mu_{k-1}^j(\theta)}{d\lambda(\theta)} \right)^{a_{ij}} \ell^i(s_k^i|\theta) d\lambda(\theta) \tag{2}$$

where $Z_k^i$ is a normalizing constant and $\lambda$ is a probability distribution on $\Theta$ with respect to which every $\mu_k^j$ is absolutely continuous. The term $d\mu_k^j(\theta)/d\lambda(\theta)$ is the Radon-Nikodym derivative of the probability distribution $\mu_k^j$. The above process starts with some initial beliefs $\mu_0^i$, $i = 1, \ldots, n$. Note that, if $\Theta$ is a finite or a countable set, the update rule in Eq. (2) reduces to: for every $B \subseteq \Theta$,

$$\mu_k^i(B) = \frac{1}{Z_k^i} \sum_{\theta \in B} \prod_{j=1}^{n} \mu_{k-1}^j(\theta)^{a_{ij}} \ell^i\left(s_k^i|\theta\right) \tag{3}$$

The updates in Eqs. (2) and (3) can be viewed as two-step processes. First, every agent constructs an aggregate belief using weighted geometric average of its own belief and the beliefs of its neighbors. Then, each agent performs a Bayes' update using the aggregated belief as a prior.

### *A. Connection with the Stochastic Mirror Descent Method*

To make this connection[2], we observe that the optimization problem in Eq. (1) is equivalent to the following problem:

$$\min_{\pi \in \mathscr{P}_\Theta} \mathbb{E}_\pi \sum_{i=1}^{n} D_{KL}\left(f^i \| \ell^i\right) = \min_{\pi \in \mathscr{P}_\Theta} \sum_{i=1}^{n} \mathbb{E}_\pi \mathbb{E}_{f^i}[-\log \ell^i]$$

where $\mathscr{P}_\Theta$ is the set of all distributions on $\Theta$. Under some technical conditions the expectations can exchange the order, so the problem in Eq. (1) is equivalent to the following one:

$$\min_{\pi \in \mathscr{P}_\Theta} \sum_{i=1}^{n} \mathbb{E}_{f^i} \mathbb{E}_\pi[-\log \ell^i] \tag{4}$$

The difficulty in evaluating the objective function in Eq. (4) lies in the fact that the distributions $f^i$ are unknown. A generic approach to solving such problems is the class of stochastic approximation methods, where the objective is minimized by constructing a sequence of gradient-based iterates where the true gradient of the objective (which is not available) is replaced with a gradient sample that is available at the given update time. A particular method that is relevant

[2]Particularly simple when $\Theta$ is a finite set.

here is the stochastic mirror-descent method which would solve the problem in Eq. (4), in a centralized fashion, by constructing a sequence $\{x_k\}$, as follows:

$$x_k = \arg\min_{y \in X} \left\{ \langle g_{k-1}, y \rangle + \frac{1}{\alpha_{k-1}} D_w(y, x_{k-1}) \right\} \tag{5}$$

where $g_k$ is a noisy realization of the gradient of the objective function in Eq. (4) and $D_w(y,x)$ is a Bregman distance function associated with a distance-generating function $w$, and $\alpha_k > 0$ is the step-size. If we take $w(t) = t \log t$ as the distance-generating function, then the corresponding Bregman distance is the Kullback-Leiblier (KL) divergence $D_{KL}$. Thus, in this case, *the update rule in Eq.* (2) *corresponds to a distributed implementation of the stochastic mirror descent algorithm* in (5), where $D_w(y,x) = D_{KL}(y,x)$ and the stepsize is fixed, i.e.., $\alpha_k = 1$ for all $k$.

We summarize the preceding discussion in the following proposition.

**Proposition 1** *The update rule in Eq.* (2) *defines a probability measure $\mu_k^i$over the set $\Theta$ generated by the probability density $\bar{\mu}_k^i = d\mu_k^i/d\lambda$ that coincides with the solution of the distributed stochastic mirror descent algorithm applied to the optimization problem in Eq.* (1)*., i.e.*

$$\bar{\mu}_k^i = \arg\min_{\pi \in P_\Theta} \left\{ \mathbb{E}_\pi[-\log \ell^i(s_k^i|\cdot)] + \sum_{j=1}^n a_{ij} D_{KL}(\pi \| \bar{\mu}_{k-1}^j) \right\} \tag{6}$$

*Proof:* We need to show that the density $\bar{\mu}_k^i$ associated with the probability measure $\mu_k^i$ defined by Eq. (2) minimizes the problem in Eq. (6).

First, define the argument in Eq. (6) as

$$G(\pi) = -\mathbb{E}_\pi \log \ell^i \left(s_k^i|\cdot\right) + \sum_{j=1}^n a_{ij} D_{KL}\left(\pi || \bar{\mu}_{k-1}^j\right)$$

and add and subtract the KL divergence between $\pi$ and the density $\bar{\mu}_k^i$ to obtain

$$\begin{aligned} G(\pi) &= -\mathbb{E}_\pi \log \ell^i(s_k^i|\cdot) + \sum_{j=1}^n a_{ij} D_{KL}(\pi \| \bar{\mu}_{k-1}^j) - D_{KL}(\pi \| \mu_k^i) + D_{KL}(\pi \| \bar{\mu}_k^i) \\ &= -\mathbb{E}_\pi \log \ell^i \left(s_k^i|\cdot\right) + D_{KL}\left(\pi \| \bar{\mu}_k^i\right) + \sum_{j=1}^n a_{ij} \mathbb{E}_\pi \left( \log \frac{\pi}{\bar{\mu}_{k-1}^j} - \log \frac{\pi}{\bar{\mu}_k^i} \right) \\ &= -\mathbb{E}_\pi \log \ell^i(s_k^i|\cdot) + D_{KL}\left(\pi \| \bar{\mu}_k^i\right) + \sum_{j=1}^n a_{ij} \mathbb{E}_\pi \log \frac{\bar{\mu}_k^i}{\bar{\mu}_{k-1}^j} \end{aligned}$$

Now, we use the relation for the density $\bar{\mu}_k^i = d\mu_k^i/d\lambda$, which is implied by the update rule for $\mu_k^i$ in Eq. (2), and obtain

$$G(\pi) = -\mathbb{E}_\pi \log \ell^i\left(s_k^i|\cdot\right) + D_{KL}\left(\pi\|\bar{\mu}_k^i\right) + \sum_{j=1}^n a_{ij}\mathbb{E}_\pi \log\left(\frac{1}{\bar{\mu}_{k-1}^j}\frac{1}{Z_k^i}\prod_{m=1}^n \left(\bar{\mu}_{k-1}^m\right)^{a_{im}} \ell^i(s_k^i|\cdot)\right)$$

$$= -\log Z_k^i + D_{KL}\left(\pi\|\bar{\mu}_k^i\right)$$

The first term in the preceding line does not depend on the distribution $\pi$. Thus, we conclude that the solution to the problem in Eq. (6) is the density $\pi^* = \bar{\mu}_k^i$ (almost everywhere). ■

## IV. Countable Hypothesis Set

In this section we present a concentration result for the update rule in Eq. (3) specific for the case of a countable hypothesis set. Later in Section V we will analyze the case of $\Theta \subset \mathbb{R}^d$.

First, we provide some definitions that will help us build the desired results. Specifically, we will study how the beliefs of all agents concentrate around the true hypothesis $\theta^*$.

**Definition 1** *Define a Kullback-Leibler Ball (KL) of radius $r$ centered at $\theta^*$ as.*

$$\mathcal{B}_r(\theta^*) = \left\{\theta \in \Theta \,\middle|\, \frac{1}{n}\sum_{i=1}^n D_{KL}\left(\ell^i\left(\cdot|\theta^*\right), \ell^i\left(\cdot|\theta\right)\right) \leq r\right\}$$

**Definition 2** *Define the covering of the set $\mathcal{B}_r^c(\theta^*)$ generated by a strictly increasing sequence $\{r_l\}_{l=1}^\infty$ with $r_1 = r$ as the union of disjoint KL bands as follows:*

$$\mathcal{B}_r^c(\theta^*) = \bigcup_{l=1}^\infty \{\mathcal{B}_{r_{l+1}}(\theta^*) \backslash \mathcal{B}_{r_l}(\theta^*)\}$$

*where $\{\mathcal{B}_{r_{l+1}} \backslash \mathcal{B}_{r_l}\}$ denotes the complement between the set $\mathcal{B}_{r_{l+1}}$ and the set $\mathcal{B}_{r_{l+1}}$, i.e. $\mathcal{B}_{r_{l+1}} \cap \mathcal{B}_{r_l}^c$ We denote the cardinality $\{\mathcal{B}_{r_{l+1}} \backslash \mathcal{B}_{r_l}\}$ by $\mathcal{N}_{r_l}$, i.e. $\left|\{\mathcal{B}_{r_{l+1}} \backslash \mathcal{B}_{r_l}\}\right| = \mathcal{N}_{r_l}$.*

We are interested in bounding the beliefs' concentration on a ball $\mathcal{B}_r\left(\theta^*\right)$ for an arbitrary $r > 0$, which is based on a covering of the complement set $\mathcal{B}_r^c\left(\theta^*\right)$. To this end, Definitions 1 and 2 provide the tools for constructing such a covering. The strategy is to analyze how the hypotheses are distributed in the space of probability distributions, see Figure 2. The next assumption will provide conditions on the hypothesis set which guarantee the concentration results.

**Assumption 3** *The series $\Sigma_{l\geq 1} \exp\left(-r_l^2 + \log \mathcal{N}_{r_l}\right)$ converges, where the sequence $\{r_l\}$ is as in Definition 2.*

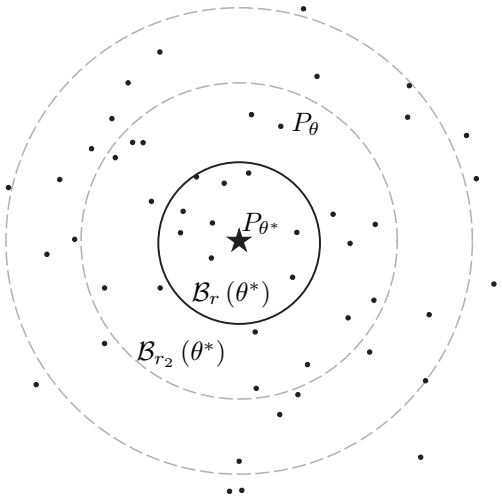


Fig. 2: Creating a covering for a ball $\mathcal{B}_r(\theta^*)$. ★ represents the correct hypothesis $\ell^i(\cdot|\theta^*)$, • indicates the location of other hypotheses and the dash lines indicates the boundary of the balls $\mathcal{B}_{r_l}(\theta^*)$.

We are now ready to state the main result for a countable hypothesis set $\Theta$.

**Theorem 1** *Let Assumptions 1, 2 and 3 hold, and let $\rho \in (0,1)$ be a desired probability tolerance. Then, the belief sequences $\{\mu_k^i\}$, $i \in V$, generated by the update rule in Eq.* (3)*, with the initial beliefs such that $\mu_0^i(\theta^*) > \epsilon$ for all $i$, have the following property: for any $\sigma \in (0,1)$ and any radius $r > 0$ with probability $1-\rho$,*

$$\mu_k^i\left(\mathcal{B}_r(\theta^*)\right) \geq 1-\sigma \qquad \text{for all } i \text{ and all } k \geq N$$

*where $N = \min_{k\geq 1}\{k \in N_1 \cap N_2\}$ with the sets $N_1$ and $N_2$ given by*

$$N_1 = \left\{k \mid \exp\left(\frac{1}{8\log^2\frac{1}{\alpha}}\right)\sum_{l\geq 1}\mathcal{N}_{r_l}\exp\left(-kr_l^2\right) \leq \rho\right\}$$

$$N_2 = \left\{k \mid C_3\exp\left(-\frac{k}{2}\gamma(\theta)\right) \leq \sigma\prod_{i=1}^n \mu_0^i(\theta)^{\frac{1}{n}}, \forall\theta \notin \mathcal{B}_r(\theta^*)\right\}$$

*where $\gamma(\theta) = \frac{1}{n}\sum_{i=1}^n D_{KL}\left(\ell^i(\cdot|\theta^*)\,||\ell^i(\cdot|\theta)\right)$, $C_3 = \frac{1}{\epsilon}\exp\left(\frac{8\log\frac{1}{\alpha}\log n}{1-\lambda}\right)$, $\alpha$ is as in Assumption 1(b), $\mathcal{N}_{r_l}$ and $r_l$ are as in Definition 2, while $\lambda = 1-\eta/4n^2$. If $A$ is a lazy-metropolis matrix, then $\lambda = 1-1/\mathcal{O}(n^2)$.*

Observe that if $k \in N_1$, then $m \in N_1$ for all $m \geq k$, and the same is true for the set $N_2$, so we can alternatively write

$$N = \max\left\{\min_{k\geq 1}\{k \in N_1\},\ \min_{k\geq 1}\{k \in N_2\}\right\}.$$

Further, note that $N$ depends on the radius $r$ of the KL ball, as the set $N_1$ involves $\mathcal{N}_{r_l}$ and $r_l$ which both depend on $r$, while the set $N_2$ explicitly involves $r$. Finally, note that the smaller the radius $r$, the larger $N$ is. We see that $N$ also depends on the number $n$ of agents, the learning parameter $\alpha$, the learning capabilities of the network represented by $\gamma(\theta)$, the initial beliefs $\mu_0^i$, the number of hypotheses that are far away from $\theta^*$ and their probability distributions.

Theorem 1 states that the beliefs of all agents will concentrate within the KL ball $\mathcal{B}_r(\theta^*)$ with a radius $r > 0$ for a large enough $k$, i.e., $k \geq N$. Note that the (large enough) index $N$ is determined as the smallest $k$ for which two relations are satisfied, namely, the relations defining the index sets $N_1$ and $N_2$. The set $N_1$ contains all indices $k$ for which a weighted sum of the total mass of the hypotheses $\theta \notin \mathcal{B}_r(\theta^*)$ is small enough (smaller than the desired probability tolerance $\rho$). Specifically, we require the number $\mathcal{N}_{r_l}$ of hypothesis in the $l$-th band does not grow faster than the squared radius $r_l^2$ of the band, i.e., the wrong hypothesis should not accumulate too fast far away from the true hypothesis $\theta^*$. Moreover, the condition in $N_1$ also prevents having an infinite number of hypothesis per band. The set $N_2$ captures the iterations $k$ at which, for all agents, the current beliefs $\mu_k^i$ had recovered from the the cumulative effect of "wrong" initial beliefs that had given probability masses to hypotheses far away from $\theta^*$.

In the proof for Theorem 1, we use the relation between the posterior beliefs and the initial beliefs on a measurable set $B$ such that $\theta^* \in B$. For such a set, we have

$$\mu_k^i(B) = \frac{1}{Z_k^i}\sum_{\theta\in B}\prod_{j=1}^{n}\mu_0^j(\theta)^{\left[A^k\right]_{ij}}\prod_{t=1}^{k}\prod_{j=1}^{n}\ell^j(s_t^j|\theta)^{\left[A^{k-t}\right]_{ij}} \tag{7}$$

where $Z_k^i$ is the appropriate normalization constant. Furthermore, after a few algebraic operations we obtain

$$\mu_k^i(B) \geq 1 - \sum_{\theta\in B^c}\prod_{j=1}^{n}\left(\frac{\mu_0^j(\theta)}{\mu_0^j(\theta^*)}\right)^{\left[A^k\right]_{ij}}\prod_{t=1}^{k}\prod_{j=1}^{n}\left(\frac{\ell^j(s_t^j|\theta)}{\ell^j(s_t^j|\theta^*)}\right)^{\left[A^{k-t}\right]_{ij}} \tag{8}$$

Moreover, since $\mu_0^j(\theta^*) > \epsilon$ for all $j$, it follows that

$$\mu_k^i(B) \geq 1 - \frac{1}{\epsilon}\sum_{\theta\in B^c}\prod_{t=1}^{k}\prod_{j=1}^{n}\left(\frac{\ell^j(s_t^j|\theta)}{\ell^j(s_t^j|\theta^*)}\right)^{\left[A^{k-t}\right]_{ij}} \tag{9}$$

Now we will state a useful result from [19] which will allow us to bound the right hand term of Eq. (8).

**Lemma 1** *[Lemma 2 in [19]] Let Assumptions 2 hold, then the matrix $A$ satisfies: for all $i$,*

$$\sum_{t=1}^{k}\sum_{j=1}^{n}\left|\left[A^{k-t}\right]_{ij}-\frac{1}{n}\right|\leq\frac{4\log n}{1-\lambda}$$

*where $\lambda=1-\eta/4n^2$, and if $A$ is a lazy-metropolis matrix associated with $\mathcal{G}$ then $\lambda=1-1/\mathcal{O}(n^2)$.*

If follows from Eq. (9), Lemma 1 and Assumption 1 that

$$\mu_k^i(B)\geq 1\text{–}C_3\sum_{\theta\in B^c}\prod_{t=1}^{k}\prod_{j=1}^{n}\left(\frac{\ell^j(s_t^j|\theta)}{\ell^j(s_t^j|\theta^*)}\right)^{\frac{1}{n}} \tag{10}$$

for all $i$, where $C_3$ is as defined in Theorem 1. Next we provide an auxiliary result about the concentration properties of the beliefs on a set $B$.

**Lemma 2** *For any $k\geq 0$ it holds that*

$$\mathbb{P}_{\boldsymbol{f}}\left(\bigcup_{\theta\in B^c}\left\{\bar{v}_k(\theta)\geq-\frac{k}{2}\gamma(\theta)\right\}\right)\leq C_2\sum_{l\geq 1}\mathcal{N}_{r_l}\exp\left(-kr_l^2\right)$$

*where $C_2=\exp\left(\frac{1}{8}\frac{1}{\log^2\frac{1}{\alpha}}\right)$, $\gamma(\theta)$ and $\mathcal{N}_{r_l}$ and $r_l$ are as in Theorem 1.*

*Proof:* First define the following random variable

$$\bar{v}_k(\theta)=\sum_{t=1}^{k}\frac{1}{n}\sum_{i=1}^{n}\log\frac{\ell^i(S_t^i|\theta)}{\ell^i(S_t^i|\theta^*)}$$

Then, by using the union bound and McDiarmid inequality we have,

$$\begin{aligned}\mathbb{P}_{\boldsymbol{f}}\left(\bigcup_{\theta\in B^c}\left\{\bar{v}_k(\theta)\text{-}\mathbb{E}_{\boldsymbol{f}}\bar{v}_k(\theta)\geq\bar{\epsilon}\right\}\right)&\leq\sum_{\theta\in B^c}\mathbb{P}_{\boldsymbol{f}}\left\{\bar{v}_k(\theta)-\mathbb{E}_{\boldsymbol{f}}\bar{v}_k(\theta)\geq\bar{\epsilon}\right\}\\&\leq\sum_{\theta\in B^c}\exp\left(\frac{2\bar{\epsilon}^2}{4k\log^2\frac{1}{\alpha}}\right)\end{aligned}$$

and by setting $\epsilon=-\frac{1}{2}\mathbb{E}_{\boldsymbol{f}}\left[\bar{v}_k(\theta)\right]$, it follows that

$$\mathbb{P}_{\boldsymbol{f}}\left(\bigcup_{\theta\in B^c}\left\{\bar{v}_k(\theta)\geq\frac{1}{2}\mathbb{E}_{\boldsymbol{f}}\left[\bar{v}_k(\theta)\right]\right\}\right)\leq\sum_{\theta\in B^c}\exp\left(\frac{\left(\mathbb{E}_{\boldsymbol{f}}\left[\bar{v}_k(\theta)\right]\right)^2}{8k\log^2\frac{1}{\alpha}}\right)$$

It can be seen that $\mathbb{E}_{\boldsymbol{f}}\left[\bar{v}_k(\theta)\right] = -k\gamma(\theta)$, thus yielding

$$\mathbb{P}_{\boldsymbol{f}}\left(\bigcup_{\theta\in B^c}\left\{\bar{v}_k(\theta) \geq -\frac{k}{2}\gamma(\theta)\right\}\right) \leq C_2 \sum_{\theta\in B^c} \exp\left(-k\gamma^2(\theta)\right)$$

Now, we let the set $B$ be the KL ball of a radius $r$ centered at $\theta^*$ and follow Definition 2 to exploit the representation of $\mathcal{B}_r^c(\theta^*)$ as the union of KL bands, for which we obtain

$$\begin{aligned}\sum_{\theta\in\mathcal{B}_r^c(\theta^*)} \exp\left(-k\gamma^2(\theta)\right) &= \sum_{l\geq 1}\sum_{\theta\in\mathcal{B}_{r_{l+1}}\setminus\mathcal{B}_{r_l}} \exp\left(-k\gamma^2(\theta)\right) \\ &\leq \sum_{l\geq 1}\mathcal{N}_{r_l}\exp\left(-kr_l^2\right)\end{aligned}$$

thus, completing the proof. ■

We are now ready to proof Theorem 1

*Theorem 1:* From Lemma 2, where we take $k$ large enough to ensure the desired probability, it follows that with probability $1-\rho$, we have: for all $k \in N_1$,

$$\begin{aligned}\mu_k^i\left(\mathcal{B}_r(\theta^*)\right) &\geq 1 - C_3 \sum_{\theta\in\mathcal{B}_r^c(\theta^*)} \exp\left(-\frac{k}{2}\gamma(\theta)\right) \\ &\geq 1 - \sum_{\theta\in\mathcal{B}_r^c(\theta^*)} \sigma\prod_{i=1}^{n}\mu_0^i\left(\theta\right)^{\frac{1}{n}} \\ &\geq 1-\sigma\end{aligned}$$

where the last inequality follows from Eq. (10) where we further take sufficiently large $k$. ■

## V. Continuum of Hypotheses

In this section we will provide the concentration results for a continuous hypothesis set $\Theta \subseteq \mathbb{R}^d$. At first, we present some definitions that we use in constructing coverings analogously to that in Section IV. In this case, however, we employ the Hellinger distance.

**Definition 3** *Define a Hellinger Ball (H) of radius $r$ centered at $\theta^*$ as.*

$$\mathcal{B}_r(\theta^*) = \left\{\theta\in\Theta \,\middle|\, \frac{1}{\sqrt{n}}h\left(\boldsymbol{\ell}\left(\cdot|\theta^*\right), \boldsymbol{\ell}\left(\cdot|\theta\right)\right) \leq r\right\}$$

**Definition 4** *Let $(M,d)$ be a metric space. A subset $S_\delta \subseteq M$ is called $\delta$-separated with $\delta > 0$ if $d(x,y) \geq \delta$ for any $x, y \in S_\delta$. Moreover, for a set $B \subseteq M$, let $\mathcal{N}_B(\delta)$ be the smallest number*

*of Hellinger balls with centers in $S_\delta$ of radius $\delta > 0$ needed to cover the set $B$, i.e., such that $B \subseteq \bigcup_{z \in S_\delta} \mathcal{B}_\delta(z)$.*

**Definition 5** *Let $\{r_l\}$ be a strictly decreasing sequence such that $r_1 = 1$ and $\lim_{l\to\infty} r_l = 0$. Define the covering of the set $\mathcal{B}_r^c(\theta^*)$ generated by the sequence $\{r_l\}$ as follows:*

$$\mathcal{B}_r^c(\theta^*) = \bigcup_{l=1}^{L_r - 1} \{\mathcal{B}_{r_l} \backslash \mathcal{B}_{r_{l+1}}\}$$

*where $L_r$ is the smallest $l$ such that $r_l \leq r$. Moreover, given a positive sequence $\{\delta_l\}$, we denote by $\mathcal{N}_{r_l}(\delta_l)$ the maximal $\delta_l$-separated subset of the set $\{\mathcal{B}_{r_l} \backslash \mathcal{B}_{r_{l+1}}\}$ and denote its cardinality by $K_l$, i.e. $K_l = |\mathcal{N}_l(\delta_l)|$. Therefore, we have the following covering of $\mathcal{B}_r^c(\theta^*)$,*

$$\mathcal{B}_r^c(\theta^*) = \bigcup_{l \geq 1}^{L_r - 1} \bigcup_{z_m \in \mathcal{N}_l(\delta_l)} \mathcal{F}_{l,m}$$

*where $\mathcal{F}_{l,m} = \mathcal{B}_{\delta_l}(z_m \in \mathcal{N}_{r_l}(\delta_l)) \cap \{\mathcal{B}_{r_l} \backslash \mathcal{B}_{r_{l+1}}\}$.*

Figure 3 depicts the elements of a covering for a set $\mathcal{B}_r^c(\theta^*)$. The cluster of circles at the top right corner represents the balls $\mathcal{B}_{\delta_l}(z_m \in \mathcal{N}_{r_l}(\delta_l))$ and for a specific case in the left of the image we illustrate the set $\mathcal{F}_{l,m}$.

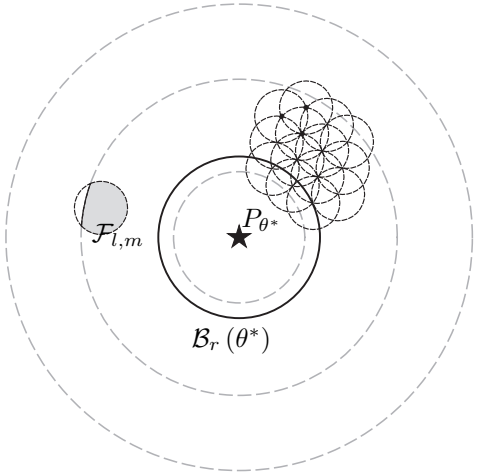


Fig. 3: Creating a covering for a set $\mathcal{B}_r(\theta^*)$. ★ represents the correct hypothesis $\ell^i(\cdot|\theta^*)$.

We are now ready to state the main result regarding continuous set of hypotheses $\Theta \subseteq \mathbb{R}^d$.

**Theorem 2** *Let Assumptions 1, 2, and 3 hold, and let $\rho \in (0,1)$ be a given probability tolerance level. Then, the beliefs $\{\mu_k^i\}$, $i \in V$, generated by the update rule in Eq.* (2) *with uniform initial beliefs, are such that, for any $\sigma \in (0,1)$ and any $r > 0$ with probability $1-\rho$,*

$$\mu_k^i\left(\mathcal{B}_r\left(\theta^*\right)\right) \geq 1-\sigma \qquad \text{for all } i \text{ and all } k \geq N$$

*where $N = \min\{k \geq 1 \,|\, k \in N_1 \text{ and } k \in N_2\}$ with*

$$N_1 = \left\{ \sum_{l=1}^{L_r-1} \exp\left( \frac{8\log\frac{1}{\alpha}\log n}{1-\lambda} - k\left(r_{l+1} - \delta_l - R\right) - d\log\delta_l \right) \leq \rho \right\}$$

$$N_2 = \left\{ \sum_{l=1}^{L_r-1} \exp\left( d\log\frac{r_l}{R} - 2k\left(r_{l+1} - \delta_l - R\right) \right) \leq \sigma \right\}$$

*for a parameter $R$ such that $r > R$ and $r_{l+1} - \delta_l - R > 0$ for all $l \geq 1$. The constant $\alpha$ is as in Assumption 1, $d$ is the dimension of the space of $\Theta$, $r_l$ and $\delta_l$ are as in Definition 5, while $\lambda$ is the same as in Theorem 1.*

Analogous to Theorem 1, Theorem 2 provides a probabilistic concentration result for the agents' beliefs around a Hellinger ball of radius $r$ with center at $\theta^*$ for sufficiently large $k$.

Similarly to the preceding section, we represent the beliefs $\mu_k^i$ in terms of the initial beliefs and the cumulative product of the weighted likelihoods received from the neighbors. In particular, analogous to Eq. (7), we have that for every $i$ and for every measurable set $B \subseteq \Theta$:

$$\mu_k^i(B) = \frac{1}{Z_k^i} \int_{\theta \in B} \prod_{t=1}^{k} \prod_{j=1}^{n} \ell^j(s_t^j|\theta)^{\left[A^{k-t}\right]_{ij}} d\mu_0^j(\theta) \tag{11}$$

with the corresponding normalization constant $Z_k^i$, and assuming all agents have uniform beliefs at time $0$.

It will be easier to work with the beliefs' densities, so we define the density of a measurable set with respect to the observed data.

**Definition 6** *The density $g_B^i$ of a measurable set $B \subseteq \Theta$, where $\mu_0^i(B) > 0$ with respect to the product distribution of the observed data is given by*

$$g_B^i(\hat{s}) = \frac{1}{\mu_0^i(B)} \int_B \prod_{t=1}^{k} \prod_{j=1}^{n} \ell^j(S_t^j|\theta)^{\left[A^{t-k}\right]_{ij}} d\mu_0^j(\theta) \tag{12}$$

*where $\hat{s} = \{S_t^i\}_{t=1:k}^{i=1:n}$ and $P_B^i = g_B^i \cdot (\lambda^i)^{\otimes k}$.*

The next lemma relates the density $g_B^i(\hat{s})$ which is defined per agent to a quantity that is common among all nodes in the network.

**Lemma 3** *Consider the densities as defined in Eq.* (12)*, then*

$$g_B^i(\hat{s}) \leq C_1 g_B(\hat{s})^{\frac{1}{n}} \tag{13}$$

*where* $C_1 = \exp\left(\frac{8\log\frac{1}{\alpha}\log n}{1-\lambda}\right)$ *and*

$$g_B(\hat{s}) = \int_B \prod_{t=1}^{k}\prod_{j=1}^{n} \ell^j(S_t^j|\theta) d\mu_0^j(\theta).$$

*Proof:* By definition of the densities, we have

$$\begin{aligned} g_B^i(\hat{s}) &= \frac{1}{\mu_0^i(B)} \int_B \prod_{t=1}^{k}\prod_{j=1}^{n} \ell^j(S_t^j|\theta)^{\left[A^{t-k}\right]_{ij}} d\mu_0^j(\theta) \\ &= \frac{1}{\mu_0^i(B)} \int_B \exp\left(\sum_{t=1}^{k}\sum_{j=1}^{n} \left[A^{t-k}\right]_{ij} \log \ell^j(S_t^j|\theta)\right) d\mu_0^j(\theta) \\ &= \frac{1}{\mu_0^i(B)} \int_B \exp\left(\sum_{t=1}^{k}\sum_{j=1}^{n} \left(\left[A^{t-k}\right]_{ij} - \frac{1}{n}\right) \log \ell^j(S_t^i|\theta) + \sum_{t=1}^{k}\frac{1}{n}\sum_{j=1}^{n} \log \ell^j(S_t^j|\theta)\right) d\mu_0^j(\theta) \end{aligned}$$

where the last line follows by adding and subtracting $1/n$. Hence, by Lemma 1, we further obtain

$$\begin{aligned} g_B^i(\hat{s}) &\leq \frac{C_1}{\mu_0^i(B)} \int_B \exp\left(\sum_{t=1}^{k}\frac{1}{n}\sum_{j=1}^{n} \log \ell^j(S_t^j|\theta)\right) d\mu_0^j(\theta) \\ &= \frac{C_1}{\mu_0^i(B)} \int_B \prod_{t=1}^{k}\prod_{j=1}^{n} \ell^j(S_t^j|\theta)^{1/n} d\mu_0^j(\theta) \\ &\leq \frac{C_1}{\mu_0^i(B)} \left(\int_B \prod_{t=1}^{k}\prod_{j=1}^{n} \ell^j(S_t^j|\theta) d\mu_0^j(\theta)\right)^{1/n} \end{aligned}$$

where the last inequality follows from Jensen's inequality. ■

The next Lemma is an analog of Lemma 2 which we use to bound the probability concentrations with respect to the ratio $g_{\mathcal{F}_{l,m}}^i(\hat{s}) / g_{\mathcal{B}_R(\theta^*)}^i(\hat{s})$.

**Lemma 4** *Consider the ratio* $g_{\mathcal{F}_{l,m}}^i(\hat{s}) / g_{\mathcal{B}_R(\theta^*)}^i(\hat{s})$*, then*

$$\mathbb{P}_{\mathcal{B}_R(\theta^*)}\left(\log\left(\frac{g_{\mathcal{F}_{l,m}}^i(\hat{s})}{g_{\mathcal{B}_R(\theta^*)}^i(\hat{s})}\right) \geq -2k(r_l - \delta_l - R)\right) \leq C_2 \exp\left(-k(r_{l+1} - \delta_l - R) + d\log\delta_l\right).$$

*with $C_2$ as defined in Lemma 2.*

*Proof:* By using the union bound, the Markov inequality and Lemma 1 in [24], we have that

$$\mathbb{P}_{\mathcal{B}_R(\theta^*)}\left(\log\left(\frac{g^i_{\mathcal{F}_{l,m}}(\hat{s})}{g^i_{\mathcal{B}_R(\theta^*)}(\hat{s})}\right)\geq y\right)\leq C_1\exp\left(-\frac{y}{2}\right)\exp\left(-k\frac{1}{n}h^2\left(\mathcal{F}_{l,m},\mathcal{B}_R(\theta^*)\right)\right)$$
$$\leq C_1\exp\left(-\frac{y}{2}-2k\left(r_{l+1}-\delta_l-R\right)\right)$$

where the last inequality follows from Proposition 5 and Corollary 1 in [24], where

$$\frac{1}{\sqrt{n}}h\left(\mathcal{F}_{l,m},\mathcal{B}_R(\theta^*)\right)\geq\frac{1}{\sqrt{n}}h\left(z_m\in\mathcal{N}_{r_l}\left(\delta_l\right),\theta^*\right)-\delta_l-R$$
$$\geq\left(r_{l+1}-\delta_l-R\right)$$

The desired result is obtained by letting $y=-2k\left(r_{l+1}-\delta_l-R\right)$. ■

Furthermore,

$$\mathbb{P}_{\mathcal{B}_R(\theta^*)}\left(\bigcup_{\mathcal{F}_{l,m}}\left\{\log\left(\frac{g^i_{\mathcal{F}_{l,m}}(\hat{s})}{g^i_{\mathcal{B}_R(\theta^*)}(\hat{s})}\right)\geq y\right\}\right)\leq\sum_{l=1}^{L_r-1}\sum_{m=1}^{K_l}\mathbb{P}_{\mathcal{B}_R(\theta^*)}\left(\log\left(\frac{g^i_{\mathcal{F}_{l,m}}(\hat{s})}{g^i_{\mathcal{B}_R(\theta^*)}(\hat{s})}\right)\geq y\right)$$
$$\leq C_2\sum_{l=1}^{L_r-1}K_l\exp\left(-k\left(r_{l+1}-\delta_l-R\right)\right)$$
$$\leq C_2\sum_{l=1}^{L_r-1}\exp\left(-k\left(r_{l+1}-\delta_l-R\right)-d\log\delta_l\right)$$

where the last inequality follows from $K_l\geq\delta_l^{-d}$, see [30], [31].

Lemma 3 allows us to represent the beliefs on the set $B^c$ as the cumulative beliefs with respect to the density $g_B(\hat{s})$. For this, similarly as in Section IV we will partition the set $B^c$ into Hellinger bands. Then for each band we will find a covering of $\delta$-separated balls and compute the concentration of probability measure with respect to density rations.

*Proof:* [Theorem 2] Lets consider the Hellinger ball $\mathcal{B}_r(\theta^*)$. We thus have

$$\mu^i_k\left(\mathcal{B}_r(\theta^*)\right)=\frac{\int\limits_{\theta\in\mathcal{B}_r(\theta^*)}\prod\limits_{t=1}^{k}\prod\limits_{j=1}^{n}\ell^j(s^j_t|\theta)^{\left[A^{k-t}\right]_{ij}}d\mu^j_0(\theta)}{\int\limits_{\theta\in\Theta}\prod\limits_{t=1}^{k}\prod\limits_{j=1}^{n}\ell^j(s^j_t|\theta)^{\left[A^{k-t}\right]_{ij}}d\mu^j_0(\theta)}$$

$$\geq 1 - \frac{\int\limits_{\theta\in\mathcal{B}_r^c(\theta^*)} \prod\limits_{t=1}^{k}\prod\limits_{j=1}^{n} \ell^j(s_t^j|\theta)^{\left[A^{k-t}\right]_{ij}} d\mu_0^j(\theta)}{\int\limits_{\theta\in\mathcal{B}_R(\theta^*)} \prod\limits_{t=1}^{k}\prod\limits_{j=1}^{n} \ell^j(s_t^j|\theta)^{\left[A^{k-t}\right]_{ij}} d\mu_0^j(\theta)}$$

$$\geq 1 - \frac{\int\limits_{\theta\in\mathcal{B}_r^c(\theta^*)} \prod\limits_{t=1}^{k}\prod\limits_{j=1}^{n} \ell^j(s_t^j|\theta)^{\left[A^{k-t}\right]_{ij}} d\mu_0^j(\theta)}{\mu_0^i\left(\mathcal{B}_R(\theta^*)\right) g^i_{\mathcal{B}_R(\theta^*)}(\hat{s})} \tag{14}$$

The construction of a partition of the set $\mathcal{B}_r^c$ presented in Definition 5 allows us to rewrite Eq. (14) as follows:

$$\mu_k^i\left(\mathcal{B}_r^c(\theta^*)\right) \leq \frac{\sum\limits_{l=1}^{L_r-1}\sum\limits_{m=1}^{K_l}\int\limits_{\mathcal{F}_{l,m}} \prod\limits_{t=1}^{k}\prod\limits_{j=1}^{n} \ell^j(s_t^j|\theta)^{\left[A^{k-t}\right]_{ij}} d\mu_0^j(\theta)}{\mu_0^i\left(\mathcal{B}_R(\theta^*)\right) g^i_{\mathcal{B}_R(\theta^*)}(\hat{s})}$$

$$= \sum_{l=1}^{L_r-1}\sum_{m=1}^{K_l} \frac{\mu_0^j(\mathcal{F}_{l,m})}{\mu_0^i\left(\mathcal{B}_R(\theta^*)\right)} \frac{g^i_{\mathcal{F}_{l,m}}(\hat{s})}{g^i_{\mathcal{B}_R(\theta^*)}(\hat{s})}$$

Finally by applying Lemma 4 and the fact that of all agents have uniform initial beliefs

$$\mu_k^i\left(\mathcal{B}_r^c(\theta^*)\right) \leq \sum_{l=1}^{L_r-1}\sum_{m=1}^{K_l} \frac{\mu_0^j(\mathcal{F}_{l,m})}{\mu_0^j\left(\mathcal{B}_R(\theta^*)\right)} \exp(-2k\left(r_{l+1}\text{-}\delta_l\text{-}R\right))$$

$$\leq \sum_{l=1}^{L_r-1} \frac{\mu_0^j\left(\mathcal{B}_{r_{l+1}}(\theta^*)\right)}{\mu_0^j\left(\mathcal{B}_R(\theta^*)\right)} \exp(-2k\left(r_{l+1}-\delta_l-R\right))$$

$$\leq \sum_{l=1}^{L_r-1} \exp\left(d\log\frac{r_{l+1}}{R} - 2k\left(r_{l+1}-\delta_l-R\right)\right)$$

The last inequality follows from the initial beliefs being uniform and the volume ratio of the two Hellinger balls with radius $r_l$ and $R$. ■

## VI. Conclusions

We proposed an algorithm for distributed learning with a countable and a continuous sets of hypotheses. Our results show non-asymptotic geometric convergence rates for the concentration of the beliefs around the true hypothesis.

While the proposed algorithm is motivated by the non-Bayesian learning models, we have shown that it is also a specific instance of a distributed stochastic mirror descent applied to a well defined optimization problem consisting of the minimization of the sum of Kullback-Liebler divergences. This indicates an interesting connection between two "separate" streams of

literature and provides an initial step to the study of distributed algorithms in a more general form. Specifically, it is interesting to explore how variations on stochastic approximation algorithms will induce new non-Bayesian update rules for more general problems. In particular, one would be interested in acceleration results for proximal methods, other Bregman distances and other constraints in the space of probability distributions.

Interaction between the agents is modeled as exchange of local probability distributions (*beliefs*) over the hypothesis set between connected nodes in a graph. This will in general generate high communication loads. Nevertheless, results are an initial study towards the distributed learning problems for general hypothesis sets. Future work will consider the effect of parametric approximation of the beliefs such that one only needs to communicate a finite number of parameters such as, for example, in Gaussian Mixture Models or Particle Filters.

## REFERENCES


[1] V. C. Gungor, B. Lu, and G. P. Hancke, "Opportunities and challenges of wireless sensor networks in smart grid," *IEEE Transactions on Industrial Electronics*, vol. 57, no. 10, pp. 3557–3564, 2010.

[2] A. Mainwaring, D. Culler, J. Polastre, R. Szewczyk, and J. Anderson, "Wireless sensor networks for habitat monitoring," in *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*. ACM, 2002, pp. 88–97.

[3] C.-Y. Chong and S. P. Kumar, "Sensor networks: evolution, opportunities, and challenges," *Proceedings of the IEEE*, vol. 91, no. 8, pp. 1247–1256, 2003.

[4] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, "Non-bayesian social learning," *Games and Economic Behavior*, vol. 76, no. 1, pp. 210–225, 2012.

[5] P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie, "Foundations of non-bayesian social learning," *Columbia Business School Research Paper*, 2015.

[6] D. Acemoglu, A. Nedić, and A. Ozdaglar, "Convergence of rule-of-thumb learning rules in social networks," in *Proceedings of the IEEE Conference on Decision and Control*, 2008, pp. 1714–1720.

[7] J. N. Tsitsiklis and M. Athans, "Convergence and asymptotic agreement in distributed decision problems," *IEEE Transactions on Automatic Control*, vol. 29, no. 1, pp. 42–50, 1984.

[8] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003.

[9] A. Nedić and A. Olshevsky, "Distributed optimization over time-varying directed graphs," *IEEE Transactions on Automatic Control*, vol. 60, no. 3, pp. 601–615, 2015.

[10] A. Olshevsky, "Linear time average consensus on fixed graphs and implications for decentralized optimization and multi-agent control," *preprint arXiv:1411.4186*, 2014.

[11] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[12] E. Mossel, A. Sly, and O. Tamuz, "Asymptotic learning on bayesian social networks," *Probability Theory and Related Fields*, vol. 158, no. 1-2, pp. 127–157, 2014.

[13] S. Shahrampour and A. Jadbabaie, "Exponentially fast parameter estimation in networks using distributed dual averagingy," in *Proceedings of the IEEE Conference on Decision and Control*, 2013, pp. 6196–6201.

[14] A. Lalitha, A. Sarwate, and T. Javidi, "Social learning and distributed hypothesis testing," in *IEEE International Symposium on Information Theory*, 2014, pp. 551–555.

[15] L. Qipeng, F. Aili, W. Lin, and W. Xiaofan, "Non-bayesian learning in social networks with time-varying weights," in *30th Chinese Control Conference (CCC)*, 2011, pp. 4768–4771.

[16] L. Qipeng, Z. Jiuhua, and W. Xiaofan, "Distributed detection via bayesian updates and consensus," in *34th Chinese Control Conference*, 2015, pp. 6992–6997.

[17] S. Shahrampour, M. Rahimian, and A. Jadbabaie, "Switching to learn," in *Proceedings of the American Control Conference*, 2015, pp. 2918–2923.

[18] M. A. Rahimian, S. Shahrampour, and A. Jadbabaie, "Learning without recall by random walks on directed graphs," *preprint arXiv:1509.04332*, 2015.

[19] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, "Distributed detection: Finite-time analysis and impact of network topology," *IEEE Transactions on Automatic Control*, vol. PP, no. 99, pp. 1–1, 2015.

[20] A. Nedić, A. Olshevsky, and C. A. Uribe, "Fast convergence rates for distributed non-bayesian learning," *preprint arXiv:1508.05161*, Aug. 2015.

[21] A. Lalitha, T. Javidi, and A. Sarwate, "Social learning and distributed hypothesis testing," *preprint arXiv:1410.4307*, 2015.

[22] A. Nedić, A. Olshevsky, and C. A. Uribe, "Network independent rates in distributed learning," *preprint arXiv: 1509.08574*, 2015.

[23] L. Su and N. H. Vaidya, "Asynchronous distributed hypothesis testing in the presence of crash failures," University of Illinois at Urbana-Champaign, Tech. Rep., 2016.

[24] L. Birgé, "About the non-asymptotic behaviour of bayes estimators," *Journal of Statistical Planning and Inference*, vol. 166, pp. 67–77, 2015.

[25] M. Rabbat and R. Nowak, "Decentralized source localization and tracking wireless sensor networks," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 3, 2004, pp. 921–924.

[26] M. Rabbat, R. Nowak, and J. Bucklew, "Robust decentralized source localization via averaging," in *IEEE International Conference on Acoustics, Speech, and Signal Processing.*, vol. 5, 2005, pp. 1057–1060.

[27] A. Nedić, A. Olshevsky, and C. A. Uribe, "Nonasymptotic convergence rates for cooperative learning over time-varying directed graphs," in *Proceedings of the American Control Conference*, 2015, pp. 5884–5889.

[28] A. Nedić, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, "On distributed averaging algorithms and quantization effects," *IEEE Transactions on Automatic Control*, vol. 54, no. 11, pp. 2506–2517, 2009.

[29] L. Birgé, "Model selection via testing: an alternative to (penalized) maximum likelihood estimators," *Annales de l'Institut Henri Poincare (B) Probability and Statistics*, vol. 42, no. 3, pp. 273 – 325, 2006.

[30] I. Dumer, "Covering spheres with spheres," *Discrete & Computational Geometry*, vol. 38, no. 4, pp. 665–679, 2007.

[31] C. Rogers, "A note on coverings," *Mathematika*, vol. 4, no. 01, pp. 1–6, 1957.